

What is claimed is:

1. A method of searching a LPM (longest prefix match) in a database which holds a plurality of prefixes in groups and defines an initial search area
5 made up of a plurality of ranges, comprising steps of:
 - (a) performing a round of binary LPM searches by executing a plurality of search instances in parallel, each search instance searching in a different range of the initial search area;
 - (b) in response to the last round of binary LPM searches, defining a new
10 search area by eliminating, from further searches, one or more ranges;
 - (c) performing a further round of binary LPM searches by executing the plurality of search instances in parallel, each search instance searching in a different sub-range of the new search area;
 - (d) in response to the last round of binary LPM searches, defining further a
15 new search area by eliminating, from further searches, one or more sub-ranges;
 - (e) storing a longest match if found in a round of binary LPM searches, and
 - (f) if necessary, repeating steps (c) to (e) to further narrow the new search area until either one of the search instances finds a longest matching prefix
20 or all the search areas have been searched, in which case the last longest match becomes the longest matching prefix.
2. The method according to claim 1 wherein the database is a routing
25 table in a packet forwarding device and the plurality of prefixes are logically sorted in groups in an ascending order of their lengths.
3. The method according to claim 2, wherein step (a) is performed with
search instances starting at predetermined locations within their respective
ranges, the predetermined locations being at about the midpoint in the lowest
30 range and being progressively shifted toward the respective low ends within higher ranges.

4. The method according to claim 3, wherein the step of defining the new search area comprises a step of:
eliminating those ranges or sub-ranges which contain prefixes shorter than the longest match of the last round of the binary LPM searches.
- 5
5. The method according to claim 4, wherein the step of performing a further round of binary LPM searches comprises further steps of:
in response to the last round of binary LPM searches, determining locations within the new search area at which the search instances start the
10 next round of binary LPM searches,
directing the search instances which searched the eliminated ranges or sub-ranges in the last round to begin the further round of binary LPM searches at the determined locations.
- 15
6. The method according to claim 5, wherein
when determining the locations within the new search area, the order of the search instances are maintained.
7. The method according to claim 3, wherein the step of executing a
20 plurality of search instances in parallel, comprises a step of:
issuing parallel memory accesses to several memory banks at once to access in parallel a plurality of bins in either the initial or new search area.
8. The method according to claim 3, wherein the step of executing a
25 plurality of search instances in parallel, comprises a step of:
issuing several memory accesses to a single memory bank to access in parallel a plurality of bins in either the initial or new search area such that the latencies of these memory accesses overlap.
- 30
9. The method according to claim 3, wherein the step of executing a plurality of search instances in parallel comprises a step of:
issuing a plurality of prefetch instructions and accessing in parallel a plurality of locations in either the initial or new search area.

10. The method according to claim 6, wherein the step of executing a plurality of search instances in parallel, comprises a step of:

5 issuing parallel memory accesses to several memory banks at once to access in parallel a plurality of bins in either the initial or new search area.

11. The method according to claim 6, wherein the step of executing a plurality of search instances in parallel, comprises a step of:

10 issuing several memory accesses to a single memory bank to access in parallel a plurality of bins in either the initial or new search area such that the latencies of these memory accesses overlap.

12. The method according to claim 6, wherein the step of executing a plurality of search instances in parallel comprises a step of:

15 issuing a plurality of prefetch instructions and accessing in parallel a plurality of locations in either the initial or new search area.

13. A method of conducting a LPM (longest prefix match) search in a packet forwarding device having a routing table containing a plurality of prefixes stored in a plurality of bins, each of which may contain one or more prefixes of the same length and markers, all the bins being logically sorted in an ascending order of their lengths and defining an initial search area which are divided into a plurality of contiguous ranges, within each of which range the bins are logically preordered for access in each round of binary LPM searches, comprising steps of:

25 (a) performing a first round of binary LPM searches by executing a plurality of search instances in parallel, each search instance searching in its respective range, starting at the bin preordered for the first access within the range;
(b) continuing further rounds of binary LPM searches by executing a plurality of search instances in parallel, each search instance searching in its
30 respective range, starting at a successively preordered bin or at one directed by a marker;

- (c) if a match or marker is found by a search instance in each round of binary LPM searches, storing it in a memory as a last longest match;
- (d) defining a new search area by eliminating, from further searches, one or more ranges containing bins of prefix lengths shorter than the last longest match;
- 5 (e) performing a further round of binary LPM searches by executing the plurality of search instances in parallel, each search instance searching in a different sub-ranges of the new search area, and
- 10 (f) if necessary, repeating steps (b) to (e) to further narrow the new search area until either one of the search instances finds a longest matching prefix or all the search areas have been searched, in which case the last longest match becomes the longest matching prefix.

14. The method according to claim 13, wherein step (e) comprises further steps of:
- 15

in response to the last round of searches, determining bins within the new search area at which the search instances start the next round of searches, and

- directing one or more search instances which searched in ranges or sub-ranges of prefixes shorter than the last longest match during the last round to begin the further round of searches starting at the determined bins of the new search area which contains the last longest match.
- 20

15. The method according to claim 14, wherein the number of bins in each ranges are predetermined and the bins ordered first for access in each range are located at about the midpoint of the lowest range and at locations progressively offset toward the low end of each of the higher ranges.
- 25

16. The method according to claim 15, wherein the step of executing a plurality of search instances in parallel, comprises a step of:
- 30

issuing parallel memory accesses to several memory banks at once to access in parallel a plurality of bins in either the initial or new search area.

17. The method according to claim 15, wherein the step of executing a plurality of search instances in parallel, comprises a step of:

5 issuing several memory accesses to a single memory bank to access in parallel a plurality of bins in either the initial or new search area such that the latencies of these memory accesses overlap.

18. The method according to claim 15, wherein the step of executing a plurality of search instances in parallel, comprises a step of:

10 issuing a plurality of prefetch instructions and accessing in parallel a plurality of bins in either the initial or new search area.

19. An apparatus for conducting LPM (longest prefix match) searches in a packet forwarding device, comprising:

15 a routing table containing a plurality of prefixes to be searched and defining an initial search area;

 a plurality of search instances for performing a plurality of rounds of parallel binary LPM searches in their respectively assigned portions of the initial search area;

20 an analyzing module for defining a new search area within the initial search area in response to the results of a last round of binary LPM searches;

 a memory for storing a longest match found in a round of binary LPM searches;

25 a controller for assigning the search instances to perform successive rounds of binary LPM searches within mutually different portions of the new search area until one of the search instances finds the longest matching prefix.

30 20. The apparatus according to claim 19, wherein the routing table comprises a plurality of bins, each of which contains one or more prefixes of a same length and may also contain at least one marker, the bins being logically sorted in order of their prefix lengths and the initial search area

being divided into a plurality of contiguous ranges, each range containing a predetermined number of bins.

21. The apparatus according to claim 20, wherein within each range, bins
5 are preordered for access by the search instances for each round of searches, if no match or marker is found.

22. The apparatus according to claim 21, wherein the ranges contain
10 sufficient number of bins to accommodate a desired number of prefixes in compliance with IPv6.

23. The apparatus according to claim 22, wherein the size of the ranges
15 are predetermined so that the worst case memory accesses are evened out across all the ranges.

24. The apparatus according to claim 23, wherein the number of bins in
each ranges are predetermined and bins ordered first for access in each
range are located at about the midpoint in the lowest range and at locations
20 progressively offset toward the low end of each of the higher ranges.

25. The apparatus according to claim 19, wherein the controller further
comprises a memory access mechanism for issuing parallel memory
accesses to several memory banks at once to access in parallel a plurality of
bins in either the initial or new search area.

26. The apparatus according to claim 19, wherein the controller further
comprises a memory access mechanism for issuing several memory
accesses to a single memory bank to access in parallel a plurality of bins in
either the initial or new search area such that the latencies of these memory
30 accesses overlap.

27. The apparatus according to claim 19, wherein the controller further
comprises a memory access mechanism for issuing a plurality of prefetch

instructions to access in parallel a plurality of bins in either the initial or new search area.

5 28. The apparatus according to claim 24, wherein the controller further comprises a memory access mechanism for issuing parallel memory accesses to several memory banks at once to access in parallel a plurality of bins in either the initial or new search area.

10 29. The apparatus according to claim 24, wherein the controller further comprises a memory access mechanism for issuing several memory accesses to a single memory bank to access in parallel a plurality of bins in either the initial or new search area such that the latencies of these memory accesses overlap.

15 30. The apparatus according to claim 24, wherein the controller further comprises a memory access mechanism for issuing a plurality of prefetch instructions to access in parallel a plurality of bins in either the initial or new search area.

20